



## **The Canonical Correlations of Color Images and their use for Demosaicing**

Yacov Hel-Or  
HP Laboratories Israel  
HPL-2003-164(R.1)  
February 23, 2004\*

color image,  
demosaic,  
image  
statistics

# The Canonical Correlations of Color Images and their use for Demosaicing

Yacov Hel-Or

Hewlett-Packard Labs

## 1 Introduction

A color image  $\mathbf{f}(x, y)$  is typically represented by three bands, each of which is a matrix of values denoting the responses of an array of photo-sensors to a captured scene. Thus, each image pixel is composed of three numbers often referred to as red (R), green (G), and blue (B) according to the spectral sensitivity of their associated sensors. Most CCD cameras, however, provide only a single value for each pixel due to their inability to position three photo-sensors at the same location. In these cases, a *mosaic* image  $m(x, y)$  is acquired, providing a single value per pixel. The sampled color values are arranged in the mosaic image according to the layout of the color filter array (CFA) in the camera's CCD. The Bayer color filter array, given in Figure 1, is a common photo-sensor arrangement in CCD arrays.

The *demosaicing* problem deals with the reconstruction of a full color image  $\mathbf{f}(x, y)$ , from its partial sampling  $m(x, y)$ . Obviously, the sampling operation is non-invertible, hence, the reconstruction problem is under-determined. In order to obtain a unique solution some prior information about the original color image  $\mathbf{f}(x, y)$  must be introduced.

Several approaches have been introduced to solve the demosaicing problem (see e.g. [23, 11, 18, 37, 36, 9, 8]). These methods assume, either explicitly or implicitly, prior knowledge about the probability of the input images  $P_F(\mathbf{f})$ . Having such a prior, the *Maximum a posteriori* (MAP) estimator can be used which chooses from among all possible solutions the one that maximizes the *a posteriori* probability  $P_{F|M}(\mathbf{f}|m)$ :

$$\hat{\mathbf{f}} = \arg \max_{\mathbf{f}} P_{F|M}(\mathbf{f}|m)$$

Using the Bayes conditional rule, this maximization can be shown to be equivalent to:

$$\hat{\mathbf{f}} = \arg \max_{\mathbf{f}} P_{M|F}(m|\mathbf{f}) P_F(\mathbf{f})$$

|   |   |   |   |   |   |
|---|---|---|---|---|---|
| R | G | R | G | R | G |
| G | B | G | B | G | B |
| R | G | R | G | R | G |
| G | B | G | B | G | B |
| R | G | R | G | R | G |
| G | B | G | B | G | B |

Figure 1: A typical arrangement of photo-sensors in a CCD array. This arrangement provides three color bands whose pixel values are partially sampled. Note that in this example, the number of green samples is twice that of the red and the blue. This corresponds to the spatial sensitivity of the human visual system at different spectral wavelengths.

The first term  $P_{M|F}$  depends solely on the noise that was introduced into the mosaic image during sampling. Since we assume that the noise characteristics are known, the first term can be easily derived. The second term,  $P_F(\mathbf{f})$ , represents the *a priori* probability of having an image  $\mathbf{f}$  as an input. Approximating the prior  $P_F(\mathbf{f})$  is a complicated task, however, it is critical and plays an important role in any reconstruction approach.

Over the last decade or so a lot of effort has been invested in an attempt to study the underlying statistics of natural images [34, 40, 30, 10, 21]. Most of this effort, however, dealt with gray-scale images, where a quite number of studies attempt to model the *spatial dependencies* existing between pixel values in natural gray-scale images. The suggested models have been used successfully in a large variety of problems, such as: image denoising, image segmentation, textures analysis and synthesis, image reconstructions, compression, perceptual grouping, and more. An excellent survey of the various models can be found in [40]. Although impressive results have been achieved by applying prior models for gray-scale images, only a few studies deal with prior models for color images. Some of these studies were developed in the context of image demosaicing, where attempts were made to characterize spatial as well as *spectral* (color) dependencies.

Demosaicing approaches that use linear shift-invariant interpolation, such as, nearest neighbor, bilinear, or bi-cubic [27, 28, 26], implicitly assume that the original image is spatially smooth. This prior stems from the inverse power law of natural images [32, 31, 14], where low frequencies are assumed to be more probable than higher frequencies. Unfortunately, this simplified model is insufficient. In addition to the fact that such a model fails to capture the spectral dependencies in color images, the smoothness assumption is far from correct

even in achromatic images especially near edges. Thus, resulting images are often smoothed out and comprise color artifacts especially adjacent to edges.

In order to preserve sharp edges, various demosaicing techniques use adaptive filtering (e.g. [11, 23, 18]), where weighted averaging schemes are applied. The weight values are typically edge-dependent and thus calculated adaptively. These techniques implicitly assume that natural images are composed of piecewise constant or piecewise linear functions [12, 25, 5]. Such techniques, although providing good reconstruction results, still do not exploit the spectral dependencies existing between color bands.

A few studies attempt to exploit spectral correlations as well as spatial correlations in order to improve reconstruction results. In a study presented by Jayanta [22], spectral dependencies are exploited through the assumption that the cross-ratio of different color bands are locally smooth. A Markov Random Field is then used to implement this characteristic in a computational process. Kimmel [23] uses the assumption that common edges are aligned along color bands. This information was derived from the assumption that a natural image forms a minimal (smooth) 2D surface in a 5-dimensional space (i.e.  $R,G,B,x,y$ ). In order to enhance reconstruction results, Kimmel suggests an iterative scheme was proposed which applies inverse diffusion along gradient directions.

Although, the latter algorithms are shown to outperform other demosaicing techniques [22], the image prior assumptions are ad-hoc and are not based on statistical inference of the image probability distribution. In the demosaicing literature, Brainard [9, 8] and Taubman [36, 37] presented solutions that were derived explicitly from statistical inference on color images. Both methods use the spatial-spectral correlations of natural images to define a prior over the image population. In both of these approaches the image prior is assumed to be a multivariate Gaussian distribution. This model of image prior was used to derive a linear estimator resulting in a non-adaptive filtering scheme.

Although these attempts to derive optimal solutions from statistical inferences are important, the underlying statistical models, as will be shown below, are not adequate and fail to capture the entire statistical dependencies that exist between color bands.

This paper suggests a demosaicing technique that is derived directly from statistical inferences on color images. Thus, the relevance of this paper is not limited to the demosaicing problem per-se, and can be used for color image de-noising, compression, segmentation, and more. The suggested technique presents a new Bayesian approach that better exploits the spectral dependencies in color images. It takes advantage of the fact that spatial discontinuities in different color bands are correlated and that this characteristic is efficiently exposed using the *Canonical Correlation Analysis* (CCA). The CCA scheme aims at finding the optimal representation of each color band such that color plane correlation is maximized [3, 19]. An introductory section about CCA, its properties, and derivations is given in the next section.

It is an interesting fact that the suggested approach which was derived solely from the statistical properties of natural images can also be derived independently from the characteristics of the human visual system. This supports the idea that the human visual system has adapted itself to the statistical properties of natural color images, and that the proposed approach is appropriate and based on a reliable statistical model.

## 2 The Canonical Correlation Analysis

### 2.1 General Definitions

The Canonical Correlation Analysis (CCA) is a technique for finding linear dependencies between two multidimensional variables. This technique which was developed by H. Hotelling in 1936 [19] aims at finding two sets of basis vectors, one for each variable, such that the correlations between the two variables in the new bases are mutually maximized. I.e., the cross correlation matrix of the two variables in the new bases is diagonal, and at the same time, the diagonal values are maximized.

In order to demonstrate the CCA operation assume two 2-dimensional variables  $\mathbf{x}$  and  $\mathbf{y}$ , whose scatter diagrams are given in Figure 2 (diagrams A and B). The two variables are linearly dependent through a linear relationship  $x_1 - x_2 = 2y_1 + y_2$ , and thus, there are two projection directions  $\mathbf{w}_x = (1, -1)^T$  and  $\mathbf{w}_y = (2, 1)^T$  in which the values  $\mathbf{w}_x^T \mathbf{x}$  and  $\mathbf{w}_y^T \mathbf{y}$  are highly correlated. These directions can be found using the CCA procedure. Diagram C in Figure 2 shows a joint scatter diagram of the two variables after they have been projected onto the directions found by the CCA procedure.




Figure 2: Joint distributions of two 2-dimensional variables: Diagram A: a scatter diagram of  $\mathbf{x}$ . Diagram B: a scatter diagram of  $\mathbf{y}$ . Diagram C: a scatter diagram of  $\mathbf{x}$  v.s.  $\mathbf{y}$  after projection onto the canonical correlation directions. Diagram D: a scatter diagram of  $\mathbf{x}$  v.s.  $\mathbf{y}$  after projection onto the direction obtained by the 1st principal component.

### 2.2 Calculating the Canonical Correlations

This section provides mathematical definitions and formulations for the Canonical Correlation scheme. This part relies heavily on [7, 6] and the interested reader is referred to these

manuscripts for further discussion.

Consider two multidimensional random variables  $\mathbf{x}$  and  $\mathbf{y}$  with zero means. The augmented covariance matrix  $C$  is defined as:

$$C = \begin{bmatrix} C_{xx} & C_{xy} \\ C_{yx} & C_{yy} \end{bmatrix} = E \left\{ \begin{pmatrix} \mathbf{x} \\ \mathbf{y} \end{pmatrix} \begin{pmatrix} \mathbf{x} \\ \mathbf{y} \end{pmatrix}^T \right\}$$

where  $E\{\cdot\}$  denotes expectation value.

Assume we are looking for two projection vectors  $\hat{\mathbf{w}}_x$  and  $\hat{\mathbf{w}}_y$  such that the correlation between  $\hat{x} = \hat{\mathbf{w}}_x^T \mathbf{x}$  and  $\hat{y} = \hat{\mathbf{w}}_y^T \mathbf{y}$  is maximized, i.e.:

$$\begin{pmatrix} \hat{\mathbf{w}}_x \\ \hat{\mathbf{w}}_y \end{pmatrix} = \arg \max_{\mathbf{w}_x, \mathbf{w}_y} \rho(\mathbf{w}_x, \mathbf{w}_y)$$

where

$$\begin{aligned} \rho &= \frac{E\{\hat{x}\hat{y}\}}{\sqrt{E\{x^2\}E\{y^2\}}} = \frac{E\{\mathbf{w}_x^T \mathbf{x} \mathbf{y}^T \mathbf{w}_y\}}{\sqrt{E\{\mathbf{w}_x^T \mathbf{x} \mathbf{x}^T \mathbf{w}_x\}E\{\mathbf{w}_y^T \mathbf{y} \mathbf{y}^T \mathbf{w}_y\}}} = \\ &= \frac{E\{\mathbf{w}_x^T C_{xy} \mathbf{w}_y\}}{\sqrt{E\{\mathbf{w}_x^T C_{xx} \mathbf{w}_x\}E\{\mathbf{w}_y^T C_{yy} \mathbf{w}_y\}}} \end{aligned}$$

It can be shown [3, 7] that the solutions for  $\hat{\mathbf{w}}_x$  and  $\hat{\mathbf{w}}_y$  satisfy the eigenvalue equations:

$$\begin{aligned} C_{xx}^{-1} C_{xy} C_{yy}^{-1} C_{yx} \hat{\mathbf{w}}_x &= \rho^2 \hat{\mathbf{w}}_x \\ C_{yy}^{-1} C_{yx} C_{xx}^{-1} C_{xy} \hat{\mathbf{w}}_y &= \rho^2 \hat{\mathbf{w}}_y \end{aligned} \quad (1)$$

Alternatively, defining matrices  $A$  and  $B$ :

$$A = \begin{bmatrix} 0 & C_{xy} \\ C_{yx} & 0 \end{bmatrix} \quad \text{and} \quad B = \begin{bmatrix} C_{xx} & 0 \\ 0 & C_{yy} \end{bmatrix}$$

the above two equations can be formulated as a single generalized eigenvalue problem:

$$A \hat{\mathbf{w}} = \rho B \hat{\mathbf{w}} \quad \text{where} \quad \hat{\mathbf{w}} = [\hat{\mathbf{w}}_x^T, \hat{\mathbf{w}}_y^T]^T$$

The set of eigenvectors satisfying the above equations are referred to as the *canonical correlation basis vectors* where the corresponding eigenvalues  $\rho^2$  are the squared *canonical correlations*. The eigenvectors  $\hat{\mathbf{w}}_x$  and  $\hat{\mathbf{w}}_y$  corresponding to the greatest eigenvalue  $\rho_1^2$  are the directions of maximal correlation value  $\rho_1$ . The subsequent 2 eigenvectors indicate the next maximal correlation directions, which are de-correlated with the previous directions, and

so on. In addition to maximizing mutual correlations, the canonical correlation basis decorrelates off-diagonal terms, i.e.  $C_{xx}$  and  $C_{yy}$  in the new basis are diagonal. It can be shown [7] that in the case where the distributions of  $\mathbf{x}$  and  $\mathbf{y}$  are Gaussians, the CCA also maximizes the mutual information between  $\mathbf{x}$  and  $\mathbf{y}$ .

There are several fundamental differences between CCA and the Principal Component Analysis (PCA). First, the CCA treats each variable as a separate entity while PCA works in a single augmented space composed of the two variable spaces together. Additionally, the CCA provides a basis which is independent of any affine transformation applied to each variable individually. Thus it takes into account only the *between* variable correlation. On the other hand, PCA provides a basis which depends also on the *within* variable correlations, and thus, might fail to provide a basis where *between* variable correlations are exposed. This behavior is illustrated in Figure 3. In this example, a 1D variable  $y$  is highly correlated with a 2D variable  $\mathbf{x}$  through a linear relationship:  $y = 1 - x_2 + n$ , where  $n$  is a noise term. Due to this correlation, all data points in the augmented space tend to be located on the plane  $x_2 + y - 1 = 0$ . As demonstrated in the figure, the *within* correlations of  $\mathbf{x}$  influence the solution of the PCA, and the 1<sup>st</sup> principal component (indicated by an arrow) does not provide useful information about the cross-variable correlation. On the other hand, the CCA provides the direction where the correlation is optimally exposed.

For the completeness of the comparison between PCA and CCA, Diagram D in Figure 2 shows the projected values of  $\mathbf{x}$  and  $\mathbf{y}$  (of diagrams A and B) onto the directions obtained by the 1<sup>st</sup> principal component. It can be seen that these directions do not provide useful information about the correlation of the two variables.

### 3 Image Prior and Color Representation

Modelling the probability distribution of gray-scale images is a challenging problem. Since typical images are composed of thousands of pixels, the probability function  $P_F(\mathbf{f})$  is defined over a very high dimensional space. Therefore, modelling the exact function is very demanding and practically impossible. Fortunately, a natural image can be viewed as a realization of a homogeneous Markov Random Field (MRF), in which the p.d.f. of a pixel value is *directly* determined by its local spatial neighborhood [15, 2, 24], and this characterization is the same over the entire image. Thus, although all pixels are mutually dependent, these dependencies are defined through successive *local* interactions, allowing the global p.d.f.  $P_F(\mathbf{f})$  to be decomposed into a product of factors. Modelling the *global* p.d.f. is then reduced to the characterization of *local* interactions between spatial neighboring pixels. In color images the local neighborhood is defined over the spatial as well as over the spectral domains. Hence, a pixel neighborhood is a 3D cube of data, composed of two spatial dimensions and one spectral dimension.

Although modelling the statistical characteristics of local spatial neighborhoods is much




Figure 3: Comparison between PCA and CCA. It can be seen that the CCA provides useful information about the linear correlation between  $\mathbf{x}$  and  $\mathbf{y}$ , while the PCA fails to expose this information.

less demanding than modelling the entire image, it is still a challenging task even for small neighborhoods such as  $3 \times 3 \times 3$ . For this reason, most of the MRF models for images use the 1st order Gauss-Markov model in which statistical dependencies are defined only between the closest neighboring pixels. One approach to handling larger neighborhoods is to approximate the neighborhood p.d.f. using marginal statistics. This is achieved by modelling the prior only over a subspace of the original neighborhood space. This subspace is defined by a set of directions onto which each image neighborhood is projected. If the dimensionality of the subspace is manageable, statistical modelling becomes practical. Clearly, marginal statistics does not necessarily provide a complete information about the global p.d.f. Therefore, it is of great importance to choose an “informative” subspace such that important information will not be lost. Using the PCA approach, the subspace with the maximal variance is chosen. This, however, does not guarantee that valuable information with respect to the specific task at hand is not lost by the projection. In our specific problem, where we are interested in inferencing on one color band given other color bands, it is preferable to choose projection directions in which the correlation between color bands is maximized. In order to achieve this the CCA approach is used.

Define a color image as a three dimensional vector field  $\mathbf{f}(x, y) = (R(x, y), G(x, y), B(x, y))^T$ . Throughout this paper we use the *log* of the linear s-RGB color values<sup>1</sup>. As an example

<sup>1</sup>The Gamma corrected values provided by digital-cameras can be used as well, as they are approximately similar to the log





Figure 4: A typical color image of a natural scene.

of a typical color image consider the image given in 4. The mutual dependencies between color values in this image can be analyzed by studying their joint histograms. Figure 5 shows three joint histograms of Red v.s. Green values, Red v.s. Blue, and Green v.s. Blue, as calculated for this image. Each of these histograms shows the joint distribution of two spectral components, and thus characterizes only the statistical dependencies along the spectral domain (neighborhood size of  $1 \times 1$ ). Figure 6 shows the full 3D scatter diagram of these RGB values. Each point in this diagram represents an RGB pixel value in the image of Figure 4. It can be seen from these histograms that the color values are correlated. Note also the line structures that appear in the scatter diagram. Table 1 shows the correlation found between the Red and the Green bands. The mutual correlations between other bands were found to be similar, and therefore were omitted from the table. The averaged correlation between the Red and the Green bands was found to be 0.91 (See Table 1) over a set of 20 768x1024 natural images.

In order to better capture the statistical dependencies between color bands we must go beyond the spectral domain and characterize the joint spatial-spectral domain. For this

---

values




Figure 5: Joint histograms of the color values in the image of Figure 4. Left to right: Red v.s. Green values, Red v.s. Blue value, and Green v.s. Blue values. The presented gray-scales are proportional to the log of the histogram values.

end we calculated the CCA directions over which the correlations between color bands are maximized. The CCA directions for the  $2 \times 1$  and  $1 \times 2$  neighborhood cases are given in Figure 7. Again, only Red and Green components are presented as the other color couples produce similar results.

It can be seen that the most correlated directions as given by the first CCA directions are the  $x$  and  $y$  spatial derivatives. The correlation values of the first CCA projections for these cases are given in Table 1, and are shown to be much higher than the pure spectral case. Figure 8 shows the joint histograms of the color values after they have been projected onto the  $x$ -derivative direction. The left histogram shows  $R_x$  v.s.  $G_x$  denoting the spatial  $x$ -derivative of the red values and the green values, respectively. Similarly, the other two histograms show  $R_x$  v.s.  $B_x$  and  $G_x$  v.s.  $B_x$ . The strong dependencies in the spatial derivatives is self evident. For comparison, Table 1 displays also the correlation values that are maintained after projecting each image neighborhood onto the first principal component using the PCA approach. It can be verified that the CCA directions produce better informative subspaces relative to those produced by the PCA.

The shape of the joint histograms, as presented in Figure 8, is typical for many natural color images. It is robust over a large variety of natural images and over all pairs of color bands. It is also robust over any directional derivative which is an outcome of the isotropic property of natural images. The characteristic distribution of these histograms suggests that directional derivatives of color images are highly probable to be located about the main diagonal  $(1, 1, 1)$  of the derivatives coordinate system. This is demonstrated in Figure 9, in which a 3D scatter diagram of  $\mathbf{f}_x(x, y) = (R_x(x, y), G_x(x, y), B_x(x, y))$  is presented. Points in this diagram represent RGB  $x$ -derivatives of pixels in the image of Figure 4. The main diagonal is plotted as a red solid line.




Figure 6: A scatter diagram of RGB color values.

The scatter diagram and the joint histograms illustrate several important aspects of the characteristics of spatial derivatives in color images: First, the spatial derivatives of different color bands, taken from the same image location, are highly probable to be similar to each other. Note, that this is not correct for the color values themselves, as demonstrated above. Second, the derivatives tend to be statistically independent if we represent the data in a coordinate system which is aligned with the main diagonal of the derivative 3D space.

We now define a new color coordinate system  $(l, c_1, c_2)$ :

$$\begin{pmatrix} l \\ c_1 \\ c_2 \end{pmatrix} = T \begin{pmatrix} R \\ G \\ B \end{pmatrix} \quad \text{where } T = \begin{pmatrix} \frac{1}{\sqrt{3}} & 0 & 0 \\ 0 & \frac{1}{\sqrt{2}} & 0 \\ 0 & 0 & \frac{1}{\sqrt{6}} \end{pmatrix} \begin{pmatrix} 1 & 1 & 1 \\ 1 & -1 & 0 \\ 1 & 1 & -2 \end{pmatrix}, \quad (2)$$

This coordinate system is widely known as the luminance/chrominance color basis (double-opponent space) where  $l$  represents the pixel luminance, and  $c_1, c_2$  the pixel chrominance [39]. This new coordinate system is a rotated version of the standard RGB coordinate system, where the luminance axis is aligned with the main diagonal, and the two chrominance axes

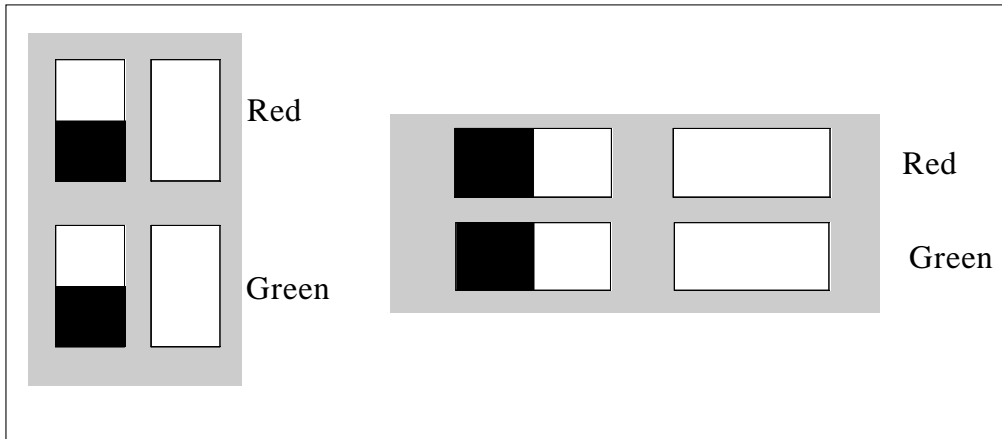


Figure 7: The Canonical Correlation basis vectors for the  $2 \times 1$  and  $1 \times 2$  cases. The upper rows correspond to the Red basis while the lower rows correspond to the Green basis. The CCA basis are ordered from the highest (left) to the lowest (right).

span the plane perpendicular to the main diagonal. The  $c_1$  axis represents the red/green component, and  $c_2$  the blue-yellow component. Since spatial derivative is a commutative operator we have:

$$T \begin{pmatrix} R_x \\ G_x \\ B_x \end{pmatrix} = T \frac{\partial}{\partial x} \begin{pmatrix} R \\ G \\ B \end{pmatrix} = \frac{\partial}{\partial x} T \begin{pmatrix} R \\ G \\ B \end{pmatrix} = \frac{\partial}{\partial x} \begin{pmatrix} l \\ c_1 \\ c_2 \end{pmatrix} = \begin{pmatrix} l_x \\ c_{1x} \\ c_{2x} \end{pmatrix}$$

thus  $(l_x, c_{1x}, c_{2x})$  is also a rotated version of the  $(R_x, G_x, B_x)$  coordinate system. As explained above, the spatial derivatives  $(l_x, c_{1x}, c_{2x})$  are, at a good approximation, statistically independent. The independence in the new basis can also be verified by calculating the mutual information which is 1.78 in the RGB-derivative basis, while only 0.06 in the new basis. These values show that the marginal probabilities in the new basis are almost independent. This observation is also in accord with [20] where applying an Independent Component Analysis (ICA) to color images yields a similar representation basis. Recall, that ICA schemes aim to find a basis whose coefficients are the least mutually dependent.

Analyzing the histogram characteristics in the new coordinate system, it can be concluded that it is highly improbable to have high derivative values in the chrominance directions. The small derivatives in the chrominance components, or alternatively, the lack of high frequencies in this domain, is demonstrated in Figure 10. It can be seen that the chrominance image tends to be smooth with relatively few discontinuities, while the luminance image is of high frequencies having a much richer collection of discontinuities.

It is interesting to relate this result to the characteristics of the human visual system. The lack of chromatic spatial sensitivity in the visual cortex [39], is in accord with the characteristics of natural color images. Thus, it may be argued that the human visual system has adapted itself to the statistical distribution of natural images.

| Subspace         | Correlation | Entropy | Mutual Inf. | Cond. Ent. |
|------------------|-------------|---------|-------------|------------|
| Pure Spectral    | 0.91        | 8.60    | 1.75        | 6.84       |
| PCA $2 \times 1$ | 0.94        | 8.19    | 1.78        | 6.40       |
| PCA $1 \times 2$ | 0.94        | 8.24    | 1.76        | 6.48       |
| PCA $3 \times 3$ | 0.94        | 8.13    | 1.86        | 6.26       |
| CCA $2 \times 1$ | 0.99        | 5.03    | 1.65        | 3.38       |
| CCA $1 \times 2$ | 0.98        | 5.04    | 1.50        | 3.54       |
| CCA $2 \times 2$ | 0.99        | 4.64    | 1.72        | 2.92       |
| CCA $3 \times 3$ | 0.99        | 4.53    | 1.68        | 2.84       |

Table 1: Statistical values for various projected subspaces. All values were calculated for the Red and Green bands, and were averaged over 20 different natural images. The statistical values are (left to right): a. The correlation between Red and Green values:  $\text{Corr}(R,G)$ . b. The differential entropy  $H(R,G)$ . c. The mutual information  $I(R,G) = H(R,G) - H(R) - H(G)$ . d. Two sided conditional entropy  $H(R|G) + H(G|R) = H(R,G) - I(R,G)$ .

## 4 Implementation: Image Demosaicing

In the following a basic statistical framework for image demosaicing is presented. Later on this framework is implemented in conjunction with the image prior model to construct an optimal solution for the demosaicing problem.

### 4.1 Image Acquisition Model

The acquisition model describes the process of acquiring an image by a digital camera. The desired image to be reconstructed is a three band image  $\mathbf{f}(x,y)$ . The image  $\mathbf{f}(x,y)$  is unavailable, however, a subset of its values is acquired by a digital camera after it has been blurred by the camera’s lens. The acquired subset is determined by the layout of the camera’s Color Filter Array (CFA). We represent the CFA arrangement by three sampling masks, which can be viewed as a 3-dimensional vector at each pixel location:

$$\mathbf{s}(x,y) = \begin{pmatrix} s^R(x,y) \\ s^G(x,y) \\ s^B(x,y) \end{pmatrix}$$

where each mask satisfies:

$$s^w(x,y) = \begin{cases} 1 & \text{if photo-sensor } w \text{ is placed at position } (x,y) \\ 0 & \text{otherwise} \end{cases} \quad w \in \{R, G, B\}$$

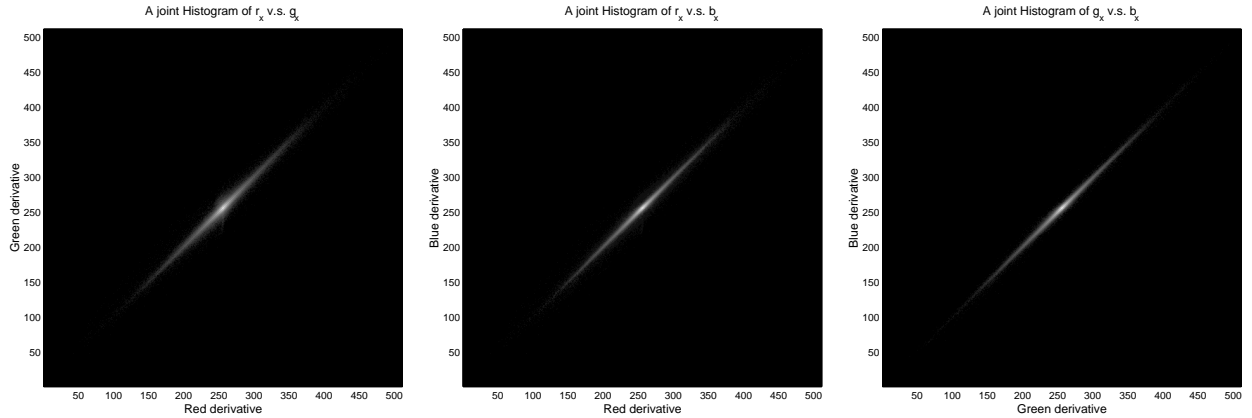


Figure 8: Joint histograms of the color x-derivatives. Left to right: x-derivative of Red v.s. x-derivative of Green, x-derivative of Red v.s. x-derivative of Blue, and x-derivative of Green v.s. x-derivative of Blue. In all histograms, gray-scales are proportional to the log of the histogram values.

When only a single color value is acquired at each pixel,  $\sum_w s^w(x, y) = 1$  for each  $(x, y)$ . The *mosaic* image  $m(x, y)$  is a result of blurring and partially sampling the original image  $\mathbf{f}(x, y)$ :

$$m(x, y) = \mathbf{s}(x, y) \cdot (H(\mathbf{f}(x, y) + \mathbf{n})) \quad (3)$$

where  $H$  is a linear operator representing the lens blur,  $\mathbf{n}$  stands for the noise incorporated during the acquisition process, and  $\cdot$  represents a vectorial inner product. The blur operator  $H$  and the statistics of the noise  $\mathbf{n}$  are assumed to be known. Since the sampling operation  $\mathbf{s}(x, y)$  is non-invertible, the solution space may include infinite many possible images. In order to obtain a unique solution, some prior information about the original color image  $\mathbf{f}(x, y)$  must be incorporated.

## 4.2 Divide and Conquer: Reconstruction vs. Enhancement

In this work we suggest a decomposition of demosaicing process into the *reconstruction* phase and the *enhancement* phase. This concept is proven to be mathematically correct when the blurring operator  $H$  is invertible. Define  $\mathbf{h}$  as the noisy and blurred image:

$$\mathbf{h} = H(\mathbf{f} + \mathbf{n})$$

Note, that  $\mathbf{h}$  is fully sampled so that  $m = \mathbf{s} \cdot \mathbf{h}$ . Using the MAP estimation, the blurred image  $\mathbf{h}$  is first estimated from  $m$ :

$$\mathbf{h}^* = \arg \max_{\mathbf{h}} P_{H|M}(\mathbf{h}|m) \quad .$$




Figure 9: A scatter diagram of the RGB derivative of the color image given in Figure 4. Each point in this diagram represents 3 values which are the RGB x-derivative of a single pixel. The red line plots a diagonal line directed toward (1, 1, 1) in the diagram coordinate system.

Since the mosaic image  $m$  is uniquely determined by  $\mathbf{h}$ , the above maximization can be rewritten:

$$\mathbf{h}^* = \arg \max_{\mathbf{h}} P_H(\mathbf{h}) \quad \text{such that} \quad m = \mathbf{s} \cdot \mathbf{h} \quad (4)$$

This optimization is referred to as the *reconstruction* phase.

Next, using  $\mathbf{h}^*$ , the enhanced image  $\mathbf{f}$  is estimated by:

$$\mathbf{f}^* = \arg \max_{\mathbf{f}} P_{F|H}(\mathbf{f}|\mathbf{h}^*)$$

This is referred to as the *enhancement* phase. The optimal solution  $\mathbf{f}^*$  using the above two phases is proven to be equivalent to the optimal solution given by a single optimization scheme (Appendix A):

$$\mathbf{f}^* = \arg \max_{\mathbf{f}} P_{F|M}(\mathbf{f}|m)$$

Since the process can be separated into two successive phases, each part can be optimized separately, enabling simple and accurate parameter settings. In our scheme, we implemented the enhancement phase using standard techniques for sharpening, such as un-sharp masking or Wiener filtering, while the reconstruction phase was implemented using a new scheme.



Figure 10: Left: The luminance component of the image presented in Figure 4. Right: One of the chrominance components ( $c_1$ )

### 4.3 Reconstruction of Color Images

In Section 3 it was demonstrated that spatial derivatives of color bands tend to be statistically independent when represented in the  $(l, c_1, c_2)$  coordinate system. Denoting a fully sampled color image

$$\mathbf{h}(x, y) = (r(x, y), g(x, y), b(x, y))^T$$

The image  $\mathbf{h}$  is represented in the new color basis:

$$\tilde{\mathbf{h}}(x, y) = T \mathbf{h}(x, y) = (l(x, y) \ c_1(x, y) \ c_2(x, y))$$

where  $T$  is defined in Equation 2. The representation in the  $(l, c_1, c_2)$  color space allows the p.d.f. of the x-derivative  $\frac{\partial}{\partial x} \tilde{\mathbf{h}}(x, y)$  to be approximated as a product of three statistically independent Gaussian marginals:

$$\begin{aligned} P \left( \frac{\partial}{\partial x} \tilde{\mathbf{h}}(x, y) \right) &= \frac{1}{Z(\alpha, \beta)} \exp \left\{ -\alpha^2 c_{1x}^2(x, y) \right\} \exp \left\{ -\alpha^2 c_{2x}^2(x, y) \right\} \exp \left\{ -\beta^2 l_x^2(x, y) \right\} = \\ &= \frac{1}{Z(\alpha, \beta)} \exp \left\{ -\alpha^2 (c_{1x}^2(x, y) + c_{2x}^2(x, y)) - \beta^2 l_x^2(x, y) \right\} \end{aligned}$$



where  $Z(\alpha, \beta)$  is a normalization factor and  $\alpha, \beta$  are constants. Figure 11 shows a joint 2D histogram of this multi-variate density function where  $\alpha = 1/5$  and  $\beta = 1/120$ . For comparison, the parametric density function is shown in the conventional  $(R, G, B)$  coordinate system. Note the similarity between the parametric density function and the histogram plotted in Figure 8. Nevertheless, the suggested density function approximates the histogram only qualitatively. A closer look at the histogram marginals suggests a super-Gaussian density function, which is characterized by a sharply peaked distribution and heavy tails (see Figure 12). Despite this, we chose to remain with the Gaussian model since it yields a linear MAP estimator while still providing good results. Using a super-Gaussian density model yields a non-linear MAP estimator, which can only be applied adaptively [33, 35, 16]. Non-linear schemes will be explored in future studies.




Figure 11: A parametric multivariate Gaussian probability function (see text) for representing color derivatives. The variance ratio between the chrominance direction and the luminance direction is 1:24.

As an outcome of the isotropic property of natural images, the  $x$  and  $y$  derivatives are statistically independent. This implies that the joint probability of both,  $x$  and  $y$  derivatives at a particular location  $(x, y)$  can be expressed by their marginal products:

$$P(\nabla \tilde{\mathbf{h}}^T(x, y)) = \frac{1}{Z(\alpha, \beta)} \exp \left\{ -\alpha^2 (\|\nabla c_1(x, y)\|^2 + \|\nabla c_2(x, y)\|^2) - \beta^2 \|\nabla l(x, y)\|^2 \right\} \quad (5)$$

where  $\nabla = \left( \frac{\partial}{\partial x}, \frac{\partial}{\partial y} \right)^T$ . Adopting the view that a natural image is a realization of a Markovian random field, and using the second-order MRF neighborhood model, a color image prior can

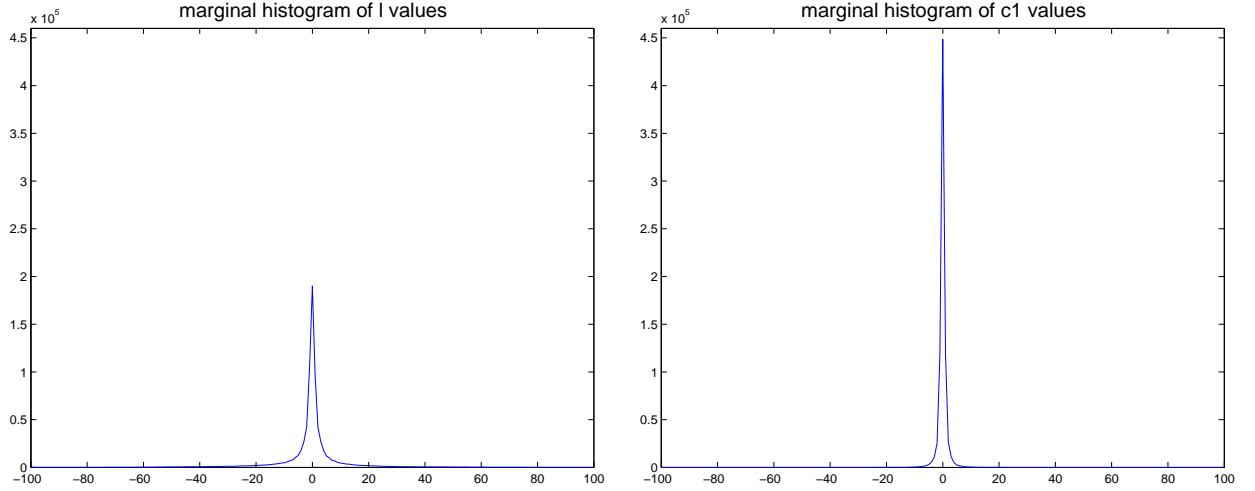


Figure 12: Plots of the marginal histograms of  $l_x$  and  $c_{1x}$  for the image given in Figure 4. Note that the marginal distributions resemble super-Gaussian density functions.

then be modelled using the following density function:

$$P_{\tilde{\mathbf{H}}}(\tilde{\mathbf{h}}) = \frac{1}{Z(\alpha, \beta)} \exp \left\{ -\alpha^2 \sum_{x,y} (\|\nabla c_1(x, y)\|^2 + \|\nabla c_2(x, y)\|^2) - \beta^2 \sum_{x,y} \|\nabla l(x, y)\|^2 \right\}$$

Referring again to the reconstruction equations (Equation 4), the (blurred) image  $\tilde{\mathbf{h}}(x, y)$  can be estimated:

$$\tilde{\mathbf{h}}^* = (l^*, c_1^*, c_2^*)^T = \arg \max_{\tilde{\mathbf{h}}} P_{\tilde{\mathbf{H}}}(\tilde{\mathbf{h}}) \quad \text{such that} \quad m = \mathbf{s} \cdot T^{-1} \tilde{\mathbf{h}}$$

Using  $P_{\tilde{\mathbf{H}}}(\tilde{\mathbf{h}})$  as defined above and taking the log of the probability function, this maximization can be expressed as a minimization problem:

$$(l^*, c_1^*, c_2^*)^T = \arg \min_{l, c_1, c_2} \left\{ \alpha^2 \sum_{x,y} (\|\nabla c_1(x, y)\|^2 + \|\nabla c_2(x, y)\|^2) + \beta^2 \sum_{x,y} \|\nabla l(x, y)\|^2 \right\} \quad (6)$$

such that  $m = \mathbf{s} \cdot T^{-1} \tilde{\mathbf{h}}$

Borrowing the terminologies from the regularization framework [13], the penalty term to be minimized is referred to as the *prior term* and the constraint equation is referred to as the *data term*. This minimization problem can be solved using the Projection onto a Convex Set (POCS) approach [4]. The approach consists of an iterative process, where at each iteration the penalty term is minimized while the constraint is imposed by projecting the intermediate results onto the constraint manifold. Since the penalty term is quadratic, the convergence is guaranteed.

Minimizing the prior term using the steepest descent method, the following iterative scheme is derived for time step  $t + 1$ :

$$\begin{aligned} l^{t+1}(x, y) &= l^t(x, y) - \beta^2 \nabla^2 l^t(x, y) \\ c_1^{t+1}(x, y) &= c_1^t(x, y) - \alpha^2 \nabla^2 c_1^t(x, y) \\ c_2^{t+1}(x, y) &= c_2^t(x, y) - \alpha^2 \nabla^2 c_2^t(x, y) \end{aligned}$$

where  $\nabla^2 = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2}$ . Applying this solution to each pixel location, the iterative scheme can be rewritten using filtering operations:

$$\begin{aligned} l^{t+1}(x, y) &= k_l(x, y) * l^t(x, y) \\ c_1^{t+1}(x, y) &= k_c(x, y) * c_1^t(x, y) \\ c_2^{t+1}(x, y) &= k_c(x, y) * c_2^t(x, y) \end{aligned}$$

where

$$\begin{aligned} k_l(x, y) &= (\delta(x, y) - \beta^2 \nabla^2) \\ k_c(x, y) &= (\delta(x, y) - \alpha^2 \nabla^2) \end{aligned}$$

and the symbol  $*$  denotes the convolution operation. There are two types of filters that are applied:  $k_l(x, y)$  which is applied to the luminance band, and  $k_c(x, y)$  which is applied to the two chrominance bands. Note, that the Laplacian  $\nabla^2$  represents a high pass filter with a radial amplitude profile  $E(u, v) = u^2 + v^2$  in the frequency domain, where  $(u, v)$  are the frequency indices. Subtracting the Laplacian filter from  $\delta(x, y)$  is in fact a low pass filter. However, since the value of  $\alpha$  is an order of magnitude larger than  $\beta$ , the two filters are quite different. Figure 13 shows the frequency profile of these two filters in the 1D frequency domain (a radial cross section). It is shown that  $k_l(x, y)$  is a “soft” low pass filter while  $k_c(x, y)$  is a “stronger” low pass filter.

Referring back to the minimization scheme, at each iteration, the *prior term* is minimized using the described filtering scheme. Following the minimization, a projection onto the measurement constraint  $m = \mathbf{s} \cdot \mathbf{h}$  must be applied. However, since the sampled values are given in the original color space, the projection is performed by transforming the intermediate result  $(l^{t+1}, c_1^{t+1}, c_2^{t+1})$  back into the  $(R, G, B)$  coordinate system:

$$\begin{pmatrix} r^{t+1} \\ g^{t+1} \\ b^{t+1} \end{pmatrix} = T^{-1} \begin{pmatrix} l^{t+1} \\ c_1^{t+1} \\ c_2^{t+1} \end{pmatrix}$$

In this original system, the values at the sampled locations are easily replace by the original mosaic image:

$$w^{t+1}(x, y) := (1 - s^w(x, y)) w^{t+1}(x, y) + s^w(x, y) m(x, y) \quad \text{for } w \in \{r, g, b\}.$$

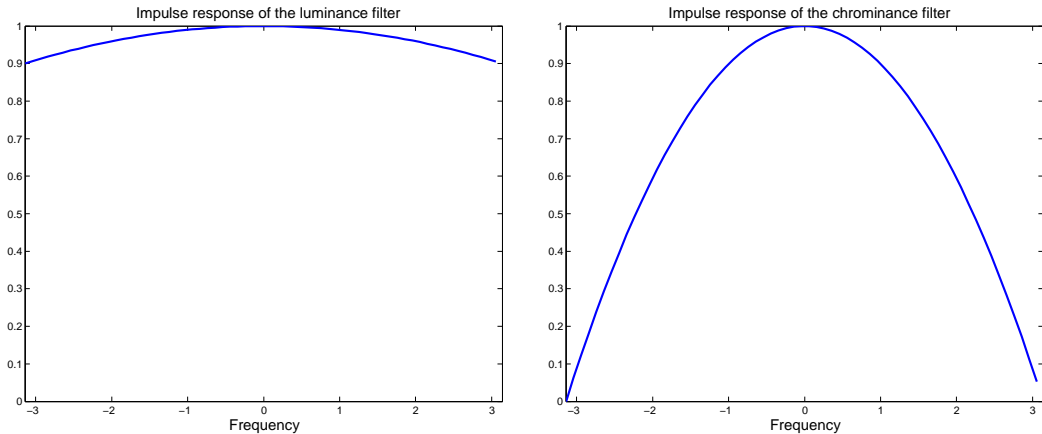


Figure 13: The frequency profile of  $k_l$  (left) compared to the frequency profile of  $k_c$  (right). Both graphs shows the radial cross section in the frequency domain.

The entire process iterates several times. Typically 4-5 iterations suffice to complete the reconstruction phase. At the end of this phase the resulting image is an optimal approximation of the blurred noisy image  $\mathbf{h}$ . The process then proceeds to the *enhancement* phase, which considers the reconstructed image  $\mathbf{h}$  as its input. The enhancement phase outputs a sharper image obtained using standard sharpening techniques, such as wiener filter or un-sharp masking. The entire algorithm is graphically illustrated in Figure 14.

An alternative approach which provides sharper results at the expense of non-linear operations can be applied as follows. Since the green band is commonly sampled more densely than the other bands (see e.g. the Bayer CFA in Figure 1), its interpolation can reproduce high frequencies which cannot be reconstructed in the other bands. Therefore, it is advantageous to treat the green band differently. In this algorithm, the green band is first interpolated. This interpolation can use interpolation techniques which are adaptive and non-linear. Among these techniques we mention the anisotropic diffusion [29], bilateral filtering [38], or robust filtering [5], all of which were shown to be equivalent [25]. Various anisotropic interpolation techniques specifically designed for the Bayer green component are suggested in the literature and can be used (see e.g. [1, 11, 23]). Following this initial step, the algorithm continues as before, with the exception that the imposed constraints are now composed of both, the original sampled values as well as the interpolated green values. Results for the two suggested algorithms are given in Section 7.

## 5 Perceptual Optimization

As mentioned above, it may be argued that the human visual system (HVS) has adapted itself to the statistical characteristics of natural images. Thus, since occurrences of high frequencies are improbable in the chrominance bands, the HSV spatial sensitivity in this domain is limited [39]. On the other hand, the HVS shows high spatial sensitivity in the

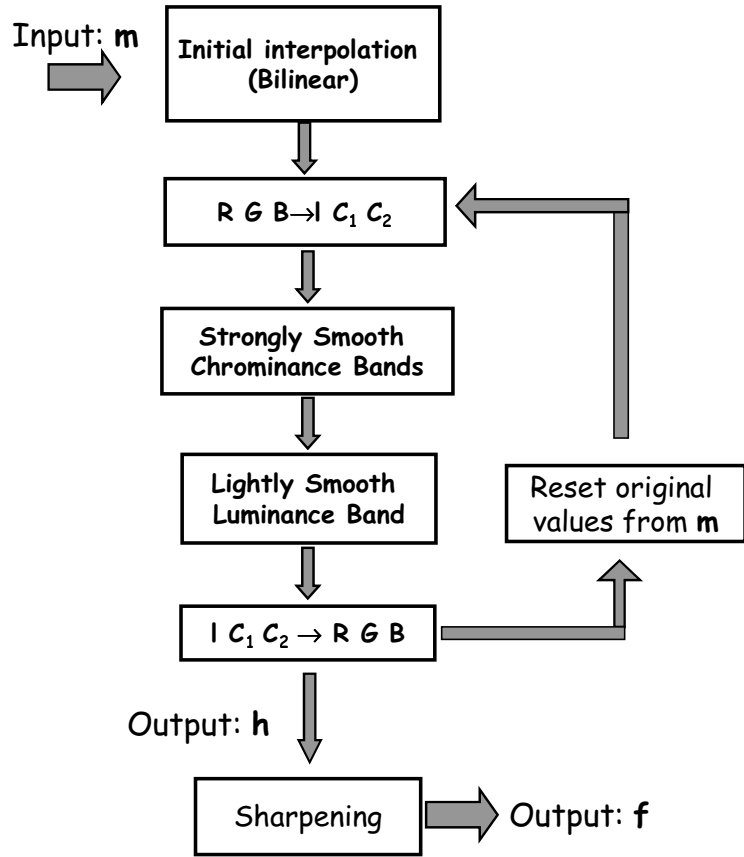


Figure 14: A block diagram describing the demosaicing algorithm.

luminance band. As a consequence, blurring the chrominance bands of a color image has a negligible effect on the perceived image (see e.g. [39]). This characteristic alone can be implemented for reducing color artifacts in demosaic images. One can simply reduce the high frequencies in the chrominance bands, smoothing out color artifacts while maintaining the luminance band unaffected. Recall, again, the minimization expression of Equation 6. This equation can be interpreted as a perceptual penalty function: The image derivatives serve as a non-smoothness measure of an image. Because natural images, in general, tend to be smooth (the inverse power spectrum law), images are penalized in this expression for being non-smooth. However, care must be taken since taking this into extreme will result in a blurred image smoothing out sharp discontinuities. Since blurred chrominance channels only slightly affect the perceived sharpness of the image, their penalty coefficient,  $\alpha$ , may be high, permitting considerable blurring in these bands. In contrast, the penalty parameter of the luminance channel is an order of magnitude smaller, and therefore defers solutions that are perceived too blurred.

## 6 Implementation

Implementation of the suggested approach is iterative. However, since each iteration is linear, it is possible to concatenate all the linear operations into a single linear operation, yielding a simple and fast solution. The concatenation of all linear operations can be easily calculated using a modification of the shift invariant property of linear systems. The input to our system is a mosaic input image, for which a shift invariant property is not valid. That is, the kernel applied to a particular pixel differs from the kernels applied to neighboring pixels. However, the same kernel is used at pixels with the same sampling configuration. For example in Figure 1, pixel locations  $(x, y)$  and  $(x + 2i, y + 2j)$  for any real  $i, j$ , have the same sampling configurations. We call this property *block shift-invariance*. In a previous study [17] a technique was developed for generating the filter kernels of block shift-invariant systems. The study describes a generalization of the impulse response of shift-invariant systems to block shift-invariant systems. The proposed technique takes as an input any computer program which applies a linear block-shift-invariant algorithm, and produces as output the equivalent set of filter kernels. These kernels can then be applied efficiently using convolution. Using this scheme, the kernels required for the suggested demosaicing process can easily be derived.

## 7 Demonstrated Results

In order to demonstrate the proposed technique, examples of results are shown along with the results of the trivial bilinear interpolation.

Figure 15 shows the original image along with its mosaic image. The mosaic image was sampled from the original image using the Bayer Color Filter Array.

Figure 16 shows three results of various demosaicing algorithms. The top image is the naive bilinear interpolation algorithm. It can be seen that many color artifacts appear in the image, mainly near edged areas. The middle image is the result of applying 5 iterations of the algorithm suggested in this paper and presented in the diagram in Figure 14. The color artifacts are drastically reduced, however the resulting image is slightly blurred relative to the original. The bottom image is a result of applying adaptive interpolation to the green component (we used bilateral filtering [38]), and applying the proposed algorithm to the red and blue components, as described above. This result presents the best performance, with almost no color artifacts, and a crisp image with sharp edges.

Using the linear block-shift-invariant properties of our algorithm, the corresponding filter kernels were extracted using the technique presented in [17]. The 12 kernels that were generated by the proposed technique, are shown in Figure 17. Each row has 4 kernels, for block positions  $(1, 1)$ ,  $(1, 2)$ ,  $(2, 1)$ , and  $(2, 2)$ , from left to right. The rows show the kernels for the Red, Green, and Blue images, from top to bottom. For comparison, Figure 18 shows



Figure 15: Top: The original RGB image. Bottom: A mosaic image which was sampled from the original image using the Bayer CFA.

the generated kernels for the bilinear algorithm.



Figure 16: Top: The resulting image after applying bilinear interpolation to the mosaiced image. Middle: The result of the proposed algorithm. Bottom: The result of applying the proposed algorithm to the red and blue components while the green component is first interpolated using adaptive filtering (see text).






Figure 17: The kernels that were calculated for the proposed demosaicing algorithm. 4 kernels are given on each row, associated with block positions (left to right)  $(1, 1)$ ,  $(1, 2)$ ,  $(2, 1)$ , and  $(2, 2)$ . The rows show the kernels for (top down) the Red, Green, and Blue images.




Figure 18: The kernels that were calculated for the bilinear algorithm. 4 kernels are given on each row, associated with block positions (left to right)  $(1, 1)$ ,  $(1, 2)$ ,  $(2, 1)$ , and  $(2, 2)$ . The rows show the kernels for (top down) the Red, Green, and Blue images.

## Acknowledgments

I would like to thank Prof. Peyman Milanfar for referring me to the literature of the Canonical Correlation Analysis.

## Appendix A

Assume  $x$  is an unknown multivariate random vector whose distribution  $P_X(x)$  is known. A measurement  $y$  is given:

$$y = SHx$$

where  $H$  and  $S$  are matrices. The MAP estimation for  $x$  yields:

$$x^* = \arg \max_x P_X(x) \quad \text{s.t.} \quad SHx = y$$

Define  $z = Hx$ , the MAP estimation for  $z$  can be calculated as well:

$$z^* = \arg \max_z P_Z(z) \quad \text{s.t.} \quad Sz = y \quad (7)$$

**Theorem 1** *If matrix  $H$  is invertible:*

$$z^* = Hx^*$$

**Proof 1:** From the the MAP definition of  $x^*$  we have

$$x^* = \arg \max_x P_X(x) \quad \text{s.t.} \quad SHx = y$$

Since  $z = Hx$  and  $H$  is invertible, we obtain:

$$Hx^* = \arg \max_z P_X(H^{-1}z) \quad \text{s.t.} \quad Sz = y$$

Multiplying the above equation by a constant will not change the maximum. Therefore,

$$Hx^* = \arg \max_z \frac{1}{|H|} P_X(H^{-1}z) \quad \text{s.t.} \quad Sz = y$$

However, since  $P_Z(z) = |H|^{-1} P_X(H^{-1}z)$ , from the above equation we obtain:

$$Hx^* = \arg \max_z P_Z(z) \quad \text{s.t.} \quad Sz = y$$

and from Equation 7 we conclude:

$$z^* = Hx^* \quad \square$$

This theorem justifies the separation of the demosaicing process into two phases, namely, the reconstruction phase, and the enhancement phase. Defining  $z$  to be the noisy blurred image  $\mathbf{h}$ , we have:  $z = H(\mathbf{f} + \mathbf{n}) = \mathbf{h}$ , and defining  $y$  to be the acquired mosaic image  $y = Sz = m$ , the reconstruction phase is associated with Equation 7. The enhancement phase deals with the estimation of  $\mathbf{f}$  given  $\mathbf{h}^*$ , which can be solved using classical techniques.

## A References

- [1] *Robust demosaicing (ROD)*, Tech. report, HP Labs, Internal Report.
- [2] R. Chellappa A. Rangarajan, *Markov random field models in image processing*, The Handbook of Brain Theory and Neural Networks (M.A. Arbib, ed.), MIT Press, 1995, pp. 564–567.
- [3] T.W. Anderson, *An introduction to multivariate statistical analysis*, John Wiley & Sons, Inc., 1958.
- [4] D.P. Bertsekas, *Nonlinear programming*, Athena Scientific, 1999.
- [5] M. Black and A. Rangarajan, *On the unification of line processes, outlier rejection, and robust statistics with applications in early vision*, International Journal of Computer Vision **19** (1996), no. 1, 57–92.
- [6] M. Borga, *Canonical correlation - a tutorial*, in <http://people.imt.liu.se/~mangus/cca/>.
- [7] ———, *Learning multidimensional signal processing*, Ph.D. thesis, Linkoping University, Department of Electrical Engineering, 1998.
- [8] D.H. Brainard, *Bayesian method for reconstructing color images from trichromatic samples*, Proc. of the IS&T /ICPS, 1994.
- [9] D.H. Brainard and D. Sherman, *Reconstructing images from trichromatic samples: From basic research to practical applications*, Proc. of the IS&T /SID, 1995.
- [10] S.C. Zhu C. E. Guo and Y. N. Wu, *Modeling visual patterns by integrating descriptive and generative models*, Int. Journal of Computer Vision **53** (2003), no. 1, 5–29.
- [11] D. R. Cok, *Reconstruction of CCD images using template matching*, Proc. IS&T's Annu. Conf. ICPS, 1994, pp. 380–385.
- [12] M. Elad, *On the bilateral filter and ways to improve it*, IEEE Transactions On Image Processing **11** (2002), no. 10, 1141–1151.

- [13] M. Jones F. Girosi and T. Poggio, *Regularization theory and neural networks architectures*, Neural Computation **7** (1995), no. 2, 219–269.
- [14] D.J. Field, *Relations between the statistics of natural images and the response profiles of cortical cells*, Journal of the Optical Society of America **4** (1987).
- [15] S. Geman and D. Geman, *Stochastic relaxation, gibbs distributions, and the bayesian restoration of images*, IEEE Trans. Pattern Anal. Mach. Intell. **6** (1984), 721–741.
- [16] M.L. Green, *Statistics of images, the TV algorithm of Rudin-Osher-Fatemi for image denoising and an improved denoising algorithm*, Tech. report, Dept. of Math., U.C.L.A.
- [17] Y. Hel-Or, *The impulse responses of block shift-invariant systems and its use for demosaicing algorithms*, Tech. Report HPL-2002-206, HP Labs (internal).
- [18] Yacov Hel-Or and Daniel Keren, *Demosaicing of color images using steerable wavelets*, Tech. Report HPL-97-104, HP Labs, August 1997.
- [19] H. Hotelling, *Relations between two sets of variables*, Biometrika **28** (1936), 321–377.
- [20] Hoyer and A. Hyvrinen, *Independent component analysis applied to feature extraction from colour and stereo images*, Network: Computation in Neural Systems (2000), no. 11(3), 191–210.
- [21] J. Huang, *Statistics of natural images and models*, Ph.D. thesis, Brown University, Applied Math. Division, May, 2000.
- [22] R. Parthasarathi Jayanta Mukherjee and S. Goyal, *Markov random field processing for color demosaicing*, Pattern Recognition Letters **22**.
- [23] R. Kimmel, *Demosaicing: Image reconstruction from color CCD samples*, EVVC (1), 1998, pp. 610–622.
- [24] S.Z. Li, *Markov random field modeling in computer vision*, Springer-Verlag, 1995.
- [25] D. Marimont M. J. Black, G. Sapiro and D. Heeger, *Robust anisotropic diffusion*, IEEE Transactions on Image Processing **7** (1998).
- [26] D. P. Mitchell and A. N. Netravali, *Reconstruction filters in computer graphics*, Computer Graphics, SIGGRAPH’88 Proceedings **22** (1988), no. 4, 221–228.
- [27] K. A. Parulski, *Color filters and processing alternatives for one-chip cameras*, IEEE Transactions on Electron Devices **32**, no. 8.
- [28] ———, *Software pixel interpolation for digital still cameras suitable for a 32-bit mcu*, IEEE Transactions on Consumer Electronics **44** (1998), no. 4, 1342–1352.

- [29] P. Perona and J. Malik, *Scale space and edge detection using anisotropic diffusion*, IEEE Trans. Pattern Anal. Mach. Intell. **12** (1990).
- [30] J. Portilla and E. P. Simoncelli, *A parametric texture model based on joint statistics of complex wavelet coefficients*, Int'l Journal of Computer Vision **40** (2000), no. 1, 49–71.
- [31] D. Ruderman, *Origins of scaling in natural images*, 1996.
- [32] D.L. Ruderman and W. Bialek, *Statistics of natural images: Scaling in the woods*, Physics Review Letters **73** (1994), no. 6, 814–817.
- [33] E. P. Simoncelli, *Bayesian denoising of visual images in the wavelet domain*, Bayesian Inference in Wavelet Based Models (P Mller and B Vidakovic, eds.), Springer-Verlag, Lecture Notes in Statistics, 1999.
- [34] ———, *Modeling the joint statistics of images in the wavelet domain*, Proc. SPIE, 44th Annual Meeting (Denver, CO), 1999, pp. 188–195.
- [35] E. P. Simoncelli and E. H. Adelson, *Noise removal via bayesian wavelet coring*, 3rd IEEE Int'l Conf on Image Processing (Laussanne, Switzerland), 1996.
- [36] D. Taubman, *Generalized wiener reconstruction of images from colour sensor data using a scale invariant prior*, Proc. of the International Conference of Image Processing, 2000.
- [37] D. Taubman, *Generalized image demosaicing and enhancement (GIDE)*, Tech. Report HPL-97-104, HP Labs (internal), Vancouver, 1996.
- [38] C. Tomasi and R. Manduchi, *Bilateral filtering for gray and color images*, IEEE International Conference on Computer Vision '98 (New Delhi, India), 1998.
- [39] B. Wandell, *Foundations of vision*, Sinaur Associates, 1995.
- [40] S. C. Zhu, *Statistical modeling and conceptualization of visual patterns*, IEEE Trans. on Pattern Analysis and Machine Intelligence (to appear).